arXiv:1412.6890v1 [stat.CO] 22 Dec 2014

Software for Distributed Computation on Medical Databases: A Demonstration Project

B. Narasimhan
Stanford University

D. L. Rubin
Stanford University

S. M. Gross
Stanford University

M. Bendersky
Stanford University

P. W. Lavori
Stanford University

Abstract

Bringing together the information latent in distributed medical databases promises to personalize medical care by enabling reliable, stable modeling of outcomes with rich feature sets (including patient characteristics and treatments received). However, there are barriers to aggregation of medical data, due to lack of standardization of ontologies, privacy concerns, proprietary attitudes toward data, and a reluctance to give up control over end use. Aggregation of data is not always necessary for model fitting. In models based on maximizing a likelihood, the computations can be distributed, with aggregation limited to the intermediate results of calculations on local data, rather than raw data. Distributed fitting is also possible for singular value decomposition. There has been work on the technical aspects of shared computation for particular applications, but little has been published on the software needed to support the "social networking" aspect of shared computing, to reduce the barriers to collaboration. We describe a set of software tools that allow the rapid assembly of a collaborative computational project, based on the flexible and extensible R statistical software and other open source packages, that can work across a heterogeneous collection of database environments, with full transparency to allow local officials concerned with privacy protections to validate the safety of the method. We describe the principles, architecture, and successful test results for the site-stratified Cox model and rank-k Singular Value Decomposition (SVD).

Keywords: distributed computation, cox regression, singular value decomposition, data aggregation, R.

1. Introduction

Bringing together the information latent in distributed medical databases promises to personalize medical care by enabling reliable, stable modeling of outcomes with rich feature sets (including patient characteristics and treatments received). To realize that promise, the National Institutes of Health (NIH) and other organizations support the aggregation of data from such databases, particularly the data developed by investigators on NIH-supported projects. These aggregated databases are made available to qualified investigators to explore and extract the useful information they contain. However, there are high (and growing) barriers to aggregation of medical data, some of them having to do with lack of standardization of

ontologies, others owing to privacy concerns, and still others arising from a generally proprietary attitude toward one's own institution's data, and a reluctance to give up control. In addition, the sheer size and complexity of some databases has caused the NIH to think about splitting the storage of aggregated databases across several centers; see for example the NIH Collaboratory (NIH 2014).

It has long been known that aggregation of data is not always necessary for model fitting. In many circumstances, such as in fitting models based on maximizing a likelihood, the computations can be distributed, with aggregation limited to the intermediate results of calculations on local data, rather than raw data (Murtagh, Demir, Jenkings, Wallace, Murtagh, Boniol, Bota, Laflamme, Boffetta, Ferretti, and Burton 2012). Indeed, sometimes distribution of the calculation among sites is necessary to share a heavy computational burden such as fitting ADMM models (Boyd, Parikh, Chu, Peleato, and Eckstein 2011). It can even help attack the "conflicting ontologies" problem, by shifting the task of translation to the sites.

There has been work on the technical aspects of shared computation (Wu, Jiang, and Ohno-Machado 2012b; Wu, Jiang, Kim, and Ohno-Machado 2012a; Jiang, Li, Wang, Wu, Xue, Ohno-Machado, and Jiang 2013) for particular applications, but little has been published on the software needed to support the "social networking" aspect of shared computing, to reduce the barriers to collaboration. We describe a set of software tools that allow the rapid assembly of a collaborative computational project, based on the flexible and extensible R (R Core Team 2014) statistical software and other open source packages, that can work across a heterogeneous collection of database environments (RedCAP, i2b2, local versions), with full transparency to allow local officials concerned with privacy protections to validate the safety of the method. The software, examples, and documentation can be found at GitHub, and it is freely modifiable by users. The rationale, scientific uses, and further details of the social networking aspects of the method are discussed elsewhere. In this article we describe the principles, architecture, and successful test results for two privacy-preserving examples. The first is a well-known distributed version of Rank-k Singular Value Decomposition and the second is a novel instance (fitting a site-stratified Cox Proportional Hazards Model). These illustrate the challenges and solutions to implementing a shared computation. We show how to take any model fit that breaks up in a similar way and implement it as a distributed computation. We describe some possible use cases that may be of interest.

2. Management of computations for a collaboration

We begin with an overview of the process and our software, with details in the Implementation section following. As illustrated in Figure 1, the setup of the software to manage computations for a collaboration begins with a few steps that need only be completed once at at a site, and that will create an environment that will support many different collaborative efforts. The local systems administrator must provide a computation server, install the application **distcomp**, and make the server URL available to the site investigators. The administrator would also set up the required permissions.

The other steps are taken by the site investigator for a given collaborative project (a "computation"). We suppose first that the investigator is the originator of a new computation and therefore, has access to a dataset that will be used as that site's contribution to the shared effort. The investigator runs a function from the **distcomp** package on his or her local com-

puter, which brings up a browser window with instructions for entering the location of the dataset, the names of its variables, and certain information about the format and values of the fields. The formula, in R format is entered. Completing this task and surviving various checks on data integrity generates a unique computation-specific identifier and a collection of metadata that define the computation and dataset.

Then the investigator sets up the slave process by running another function in the **distcomp** package. The result is that a copy of the local dataset and the metadata generated in the previous step is placed in the computation server under the unique computation identifier. At this point the originating investigator is ready to recruit other collaborators. They will have to take the step of creating a computation server, as described above, unless it has already been done at that site. Then, they will run the same setup function locally, and enter the computation definition (metadata) sent by the originating investigator. The function will return a URL (pointing to the computation server) that is to be transmitted back to the originating investigator.

Once the originating investigator has the URLs from the sites in hand, he or she runs the function for setting up a master site, and then finally another function that runs the computation. The process of setting up slave and master sites can be repeated for any number of different computations and datasets.

We imagine that a group of collaborators interested in a particular topic (say, prediction of outcome of therapy for a subtype of melanoma) would run (and subsequentially update) several such computations, perhaps variations of models. The distributed computations of interest to that group investigators might be listed and managed on a website where a registry of active and proposed projects is maintained. Depending on the scope of the collaboration there may be several such websites. One may imagine a website providing the following capabilities for a collaboration:

Registration Registering users, providing credentials, bonafides etc.

Available Methods A mechanism for adding to a growing list of statistical computations that may be implemented in the **distcomp** projects. The latest version of the package will automatically provide an up-to-date list.

Project Linkages Associating users with projects so that they may participate in the collaboration. In order to prevent malicious use and avoid inappropriate computational burdens on already established projects there are some checks and balances, such as requiring a new user to bring new data to the collaboration.

Project Examples/Templates Providing some testable examples for potential collaborators so that there is a complete understanding of the requirements of each site. Some canonical templates will be provided. Users may run them on local sites to completely examine the entire process.

Project Publishing Creation of projects and publishing them so that they move into place on a publicly visible website. A group of users or a single user will create a new collaborative project and lead it.

Tools Downloadable tools, virtual machines and detailed tool documentation for enabling dryrun experiments locally.

Documentation Wikis for facilitating communication within the collaboration.

There are many open source tools that can be used for some of the tasks described above. We focus our attention on those tasks for which new tools required development. These are the tools that allow rapid setup of a computation, at both master and slave sites, as described above, with minimal ongoing burden once the initial steps are taken. The intent is to lower the bar to collaboration to the point where it is as easy to set up a distributed collaboration as it would be to even begin the discussion of data aggregation.

3. Implementation

Our implementation is in the form of the **distcomp** R package which utilizes other R packages notably, **opencpu** and **shiny** (RStudio and Inc. 2014), to provide the infrastructure support. We use a master-slave model, where the master is the one in charge of the overall computation (the main iteration in an optimization for example) and the slaves merely compute functions on local datasets and return the function result. Both the master and slave sites are expected to have our package **distcomp** installed (from CRAN). In the course of a single fit, the master process will make an unspecified number of computation requests to the slave sites over secure `http` protocol.

The **distcomp** package is designed to be used in conjunction with an OpenCPU (Ooms 2014) server or an Rserve (Urbanek 2013) process. In fact, several other analogs may be used, but our current implementation is explicitly geared to **opencpu** because it is straightforward to install and configure. Therefore, in what follows, we assume an **opencpu** R server.

We summarize the main steps in implementing a distributed computation, followed by details of each step. Figure 1 illustrates the first three steps.

Define the Computation An R function `defineNewComputation` in **distcomp** lets one define the formula and variables to be used in the Cox model. The result of this step is an output file, the 'computation definition' or *compdef* for brevity, that will encode the characteristics of the computation. The file includes the R model formula and variables, but no data and may therefore be shared freely. The *compdef* file is used in all the ensuing steps.

Set up a Slave Process for the Computation The function `setupSlave` uses the *compdef* to create a slave process that can participate in the computation. This requires an R Server process that is suitably configured for talking across networks. Several checks are built into this step and if everything is successful, the server becomes ready to perform the new computation using the provided *compdef*.

Build a Master Process for the Computation The function `setupMaster` allows a user to create a master process (R code) for the computation that will run the computation and provide the final results of a fit, such as the coefficients and estimates of variance. It takes as input the *compdef* created in the previous step, URLs of sites participating in the computation and outputs an R script which may be run.

Fit the Model The master process code can now be run to to do the model fitting.

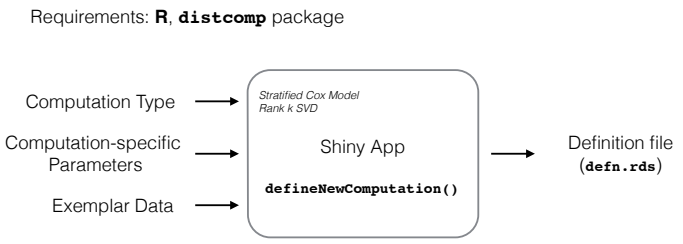

(a) Defining the computation. A computation may be define on any machine where R and the **distcomp** package are installed. The function `defineNewComputation()` launches a **shiny** app that leads the user through the process. The end result is an R data file unambiguously defining the computation instance for **distcomp**

(b) Setting up the slave. This requires a one-time configuration of an **opencpu** server with the R package **distcomp** and a writable workspace. All interaction is through a **shiny** app that configures the slave for the computation.

(c) Setting up the master. The **shiny** app accepts the URLs of all slave sites, checks them and writes an R script for the computation.

Figure 1: Inputs, outputs, and associated R functions in package **distcomp** for each task.

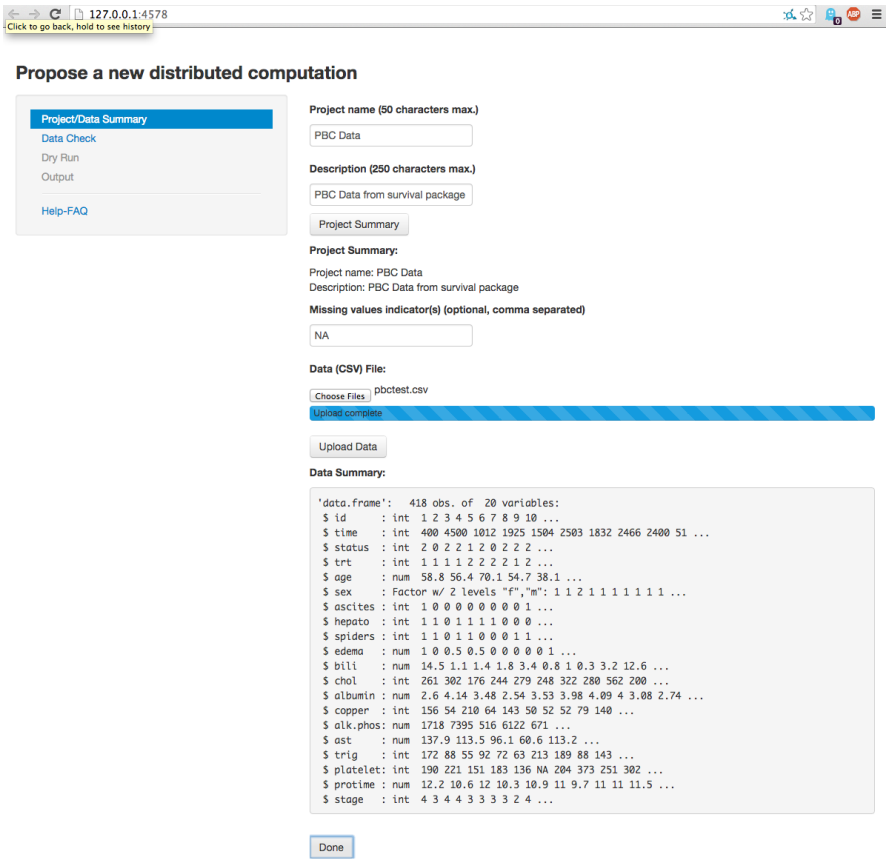
127.0.0.1:4578
Click to go back, hold to see history
Propose a new distributed computation
Project/Data Summary
Data Check
Dry Run
Output
Help-FAQ
Project name (50 characters max.)
PBC Data
Description (250 characters max.)
PBC Data from survival package
Project Summary
Project Summary:
Project name: PBC Data
Description: PBC Data from survival package
Missing values indicator(s) (optional, comma separated)
NA
Data (CSV) File:
Choose Files
pbctest.csv
Upload complete
Upload Data
Data Summary:
'data.frame': 418 obs. of 20 variables:
$ id : int 1 2 3 4 5 6 7 8 9 10 ...
$ time : int 400 4500 1012 1925 1504 2503 1832 2466 2400 51 ...
$ status : int 2 0 2 2 1 2 0 2 2 2 ...
$ trt : int 1 1 1 1 2 2 2 1 1 2 ...
$ age : num 58.8 56.4 70.1 54.7 38.1 ...
$ sex : Factor w/ 2 levels "f","m": 1 1 2 1 1 1 1 1 1 1 ...
$ ascites : int 1 0 0 0 0 0 0 0 0 1 ...
$ hepato : int 1 1 0 1 1 1 1 0 0 0 ...
$ spiders : int 1 1 0 1 1 0 0 0 1 1 ...
$ edema : num 1 0 0.5 0.5 0 0 0 0 0 1 ...
$ bili : num 14.5 1.1 1.4 1.8 3.4 0.8 1 0.3 3.2 12.6 ...
$ chol : int 261 302 176 244 279 248 322 280 562 200 ...
$ albumin : num 2.6 4.14 3.48 2.54 3.53 3.98 4.09 4 3.08 2.74 ...
$ copper : int 156 54 210 64 143 50 52 52 79 140 ...
$ alk.phos: num 1718 7395 516 6122 671 ...
$ ast : num 137.9 113.5 96.1 60.6 113.2 ...
$ trig : int 172 88 55 92 72 63 213 189 88 143 ...
$ platelet: int 190 221 151 183 136 NA 204 373 251 302 ...
$ protime : num 12.2 10.6 12 10.3 10.9 11 9.7 11 11 11.5 ...
$ stage : int 4 3 4 4 3 3 3 3 2 4 ...
Done

Figure 2: Defining a New Computation

The master and slave processes will communicate via simple *messages*. This requires a robust serialization protocol since the R data structures (consisting of vectors, matrices and lists) will be passing over the wire. The current implementation uses the simple JSON serialization format for transmitting the data structures as it is the best supported format in **opencpu**.

A typical computation will involve the optimization of some criterion, very often a likelihood function. For any specific computation, the associated master will initiate this optimization.

3.1. Define a Computation

The function `defineNewComputation` invokes a **shiny** web application and Figure 1 gives an overview of what is expected by the application. A sample screen shot is shown in Figure 2. Besides gathering some generic information such as a name and title for the computation, a formula for the model and a starting dataset is also taken as input. We expect that the initial proposer has actual data on subjects from her site in some analyzable form (CSV assumed here) with meaningful names for the covariates. Other screens enable the user to choose from the available computations, check data for conformability and execute a dry run. For example, the provided formula is checked against the dataset to ensure the model can actually be fit. The process is not complete until the fit proceeds without error; the various buttons in the user interface are articulated appropriately. The variables in the formula then become the variables that *all* other participating sites will have to provide at a minimum.

An identifier (for all practical purposes unique) iqs generated and associated with every computation definition. This ensures that a slave site (i.e. the associated R computation engine) may participate in more than one computation; the identifier serves to unambiguously distinguish various computations and associate appropriate data and functions for that computation.

The final result is a *compdef* object, containing all necessary metadata defining the computation (but no individual patient data), that is saved in an `.rds` file.

3.2. Build Slave Process for the Computation

Before a slave site can be configured for a computation, an **opencpu** server needs to be set up and its profile modified so that the **distcomp** package is loaded and configured with a workspace. The workspace configuration requires a one-time editing of a security policy (at least on the commonly used Ubuntu Linux servers) so that the **opencpu** server may write serialized objects into the area.

The function `setupSlave` invokes another **shiny** webapp that requires several inputs: the URL of the **opencpu** server, the *compdef* (see section 3.1) and the site specific data. The shiny webapp performs sanity checks, ensures that the model can be fit and if everything succeeds, it uploads information to the **opencpu** server so that the computation becomes accessible. The site is then ready to participate in a computation.

3.3. Build Master Process for the Computation

The function `setupMaster` invokes a shiny webapp that creates a master process R object that can interact with slave processes. Inputs to this application are the the *compdef* generated in section 3.1 and URLs of sites participating in the computation. Once again some sanity checks are performed to ensure that sites are indeed addressable and finally, the R code for running

the computation is generated. The application updates the user R profile with code that ensures the instantiation of an object with the correct identifier and variables and formulae. No data is needed at this point, but the master object can do no real computation until slave sites are added.

4. Two Examples

We present two self-contained examples that may be run in an R session using the **distcomp** package in conjunction with the **opencpu** R package. The first example is more detailed to show R code behind the **shiny** applications in Figure 1.

The **distcomp** package may be installed from Github as follows.

```
library("devtools")
install_github("hrpcisd/distcomp")
```

Since we use a single **opencpu** server on a single machine, the data for each site has to be kept separate and the implementation ensures that the server performs computations only with site-specific data. Furthermore, an empty directory has to be set aside for the workspace for the **opencpu** server. This is best done by modifying the user R profile to contain the following two lines.

```
library("distcomp")
distcompSetup(workspace = "full path of workspace",
              ssl.verifyhost=FALSE,
              ssl.verifypeer=FALSE)
```

(On Unix and MacOS, the user's `.Rprofile` file in the home directory will suffice, but on Windows, this needs to be inserted into the site profile.) Thus every R invocation will have loaded the **distcomp** package and know about the workspace.

The full structure of the workspace is shown in Figure 3. The `defn` folders store *compdef* objects under various identifiers generated for the computations. The `instances` folders store the instantiated objects used in the computations; these may be repeatedly saved during iterations as they change state.

4.1. A Rank-k Approximation Example

We consider the problem of approximating a matrix by another low-rank matrix. Assuming that the original matrix $\boldsymbol{X}$ is row-partioned into sub-matrices $\boldsymbol{X}_j$ at $j = 1 \ldots, K$ sites, there is a well-known iterative singular value decomposition algorithm (see appendix B) to obtain a low-rank approximation using the singular vectors, which is implemented in **distcomp**. The example below assumes three sites.

```
> print(availableComputations())
$StratifiedCoxModel
[1] "Stratified Cox Model"

$RankKSVD
[1] "Rank K SVD"
```

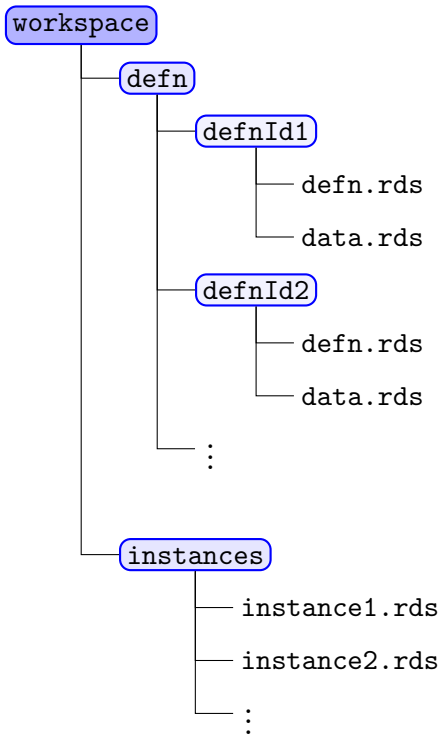

Figure 3: Structure of Workspace Area

Using this information, we construct the *compdef* for this computation and use a generic identifier `SVD`.

```
svdDef <- list(compType = names(availableComputations())[2],
               defnId = "SVD")
```

We now generate random data for three sites, start the **opencpu** server, and set the URLs for the sites to be the (same) local **opencpu** URL.

```
> set.seed(12345)
> nSites <- 3
> siteNames <- sapply(seq.int(nSites),
 +   function(i) paste("site", i, sep=""))
> siteData <- lapply(seq.int(nSites),
+    function(i) matrix(rnorm(100), nrow=20))
> library("opencpu")
  ...
  OpenCPU started.
  ...
> siteURLs <- lapply(seq.int(nSites), function(i) opencpu$url())
> names(siteData) <- names(siteURLs) <- siteNames
```

Next, we upload the data to the three sites, ensuring that we use different names for the data files for each site.

```
> siteDataFiles <- lapply(seq.int(nSites),
+   function(i) paste("site", i, ".rds", sep=""))
> ok <- Map(uploadNewComputation, siteURLs,
+   lapply(seq.int(nSites), function(i) svdDef),
+   siteData,
+   siteDataFiles)
> stopifnot(all(ok))
```

At this point, the sites are instantiated and ready to compute. So we instantiate the master, specifying that we are using a local server, which ensures that our **opencpu** server handles data for each sites appropriately.

```
> master <- svdMaster$new(defnId = svdDef$defnId, localServer=TRUE)
> for (i in seq.int(nSites)) {
+     master$addSite(siteNames[i], siteURLs[[i]],
+        dataFileName=siteDataFiles[[i]])
+ }
> result <- master$run(k=1)
> result
$v
            [,1]
[1,]  0.17947030
[2,]  0.78969198
[3,] -0.21294972
[4,] -0.54501407
[5,]  0.04229739

$d
[1] 9.707451
```

This returns the approximate first singular value and the associated vector (up to a sign change). All five singular vectors can be obtained as shown below. We also compare the results to the SVD on the aggregated matrix.

```
> result <- master$run(k=5, thr=1e-14, max.iter=100)
> x <- do.call(rbind, siteData)
> svd(x)$d
[1] 9.707537 8.199827 7.982888 7.257286 6.235182
> result$d
[1] 9.707451 8.200043 7.982650 7.257355 6.235351
> svd(x)$v
            [,1]        [,2]        [,3]        [,4]        [,5]
[1,] -0.17946375  0.08268613 -0.01644895 -0.98010572 -0.00883063
[2,] -0.78963831  0.34694371  0.34328503  0.16509457  0.33316749
[3,]  0.21305901  0.91839439 -0.25083926  0.04461477 -0.21505068
[4,]  0.54504905  0.16843629  0.53318714 -0.10009622  0.61663844
```

```
[5,] -0.04232602 -0.03120945 -0.73121540  0.01126215  0.68002329
> result$v
            [,1]        [,2]       [,3]        [,4]         [,5]
[1,]  0.17947030  0.08275684  0.0165604  0.98008722 -0.008933396
[2,]  0.78969198  0.34634459 -0.3437723 -0.16504730  0.333181988
[3,] -0.21294972  0.91875219  0.2496210 -0.04479619 -0.214978886
[4,] -0.54501407  0.16784298 -0.5334277  0.10025749  0.616612820
[5,]  0.04229739 -0.03032954  0.7312254 -0.01140918  0.680060781
```

As we can see, the distributed SVD is able to recover the same factors (up to sign change) as the standard SVD algorithm. There is a slight loss of numerical precision in the alternating power algorithm versus the standard LAPACK implementation. The JSON serialization format employed here also loses precision. Future implementations will use a more portable format such as Google protocol buffers, as we note in section 6.

5. A Distributed Stratified Cox Model Example

We next present a distributed Statified Cox model fit using the UIS dataset from Hosmer and Lemeshow (Hosmer, Lemeshow, and May 2008) on time until return to drug use for patients enrolled in two different residential treatment programs. Assuming all data in one place, one would fit a stratified cox proportional hazard model using `site` (0 or 1) as a stratification variable.

```
> uis <- readRDS("uis.RDS")
> coxOrig <- coxph(formula=Surv(time, censor) ~ age + becktota +
+                       ndrugfp1 + ndrugfp2 + ivhx3 +
+                       race + treat + strata(site), data=uis)
> summary(coxOrig)
Call:
coxph(formula = Surv(time, censor) ~ age + becktota + ndrugfp1 +
ndrugfp2 + ivhx3 + race + treat + strata(site), data = uis)

n= 575, number of events= 464
(53 observations deleted due to missingness)

               coef exp(coef)  se(coef)      z Pr(>|z|)
age       -0.028076  0.972315  0.008131 -3.453 0.000554 ***
becktota   0.009146  1.009187  0.004991  1.832 0.066914 .
ndrugfp1  -0.521973  0.593349  0.124424 -4.195 2.73e-05 ***
ndrugfp2  -0.194178  0.823512  0.048252 -4.024 5.72e-05 ***
ivhx3TRUE  0.263634  1.301652  0.108243  2.436 0.014868 *
race      -0.240021  0.786611  0.115632 -2.076 0.037920 *
treat     -0.212616  0.808466  0.093747 -2.268 0.023331 *
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
...
```

To get the same results using a distributed computation, we proceed as before, executing the steps in order.

```
> coxDef <- list(compType = names(availableComputations())[1],
+  formula = paste("Surv(time,censor)~",
       "age+becktota+ndrugfp1+ndrugfp2+ivhx3+race+treat"),
+  defnId = "UIS")
> siteData <- with(uis, split(x=uis, f=site))
> nSites <- length(siteData)
> siteNames <- sapply(seq.int(nSites),
+   function(i) paste("site", i, sep=""))
> siteURLs <- lapply(seq.int(nSites), function(i) opencpu$url())
> names(siteData) <- names(siteURLs) <- siteNames
> siteDataFiles <- lapply(seq.int(nSites),
+   function(i) paste("site", i, ".rds", sep=""))
> ok <- Map(uploadNewComputation, siteURLs,
+           lapply(seq.int(nSites), function(i) coxDef),
+           siteData,
+           siteDataFiles)
stopifnot(all(ok))
> master <- coxMaster$new(defnId = coxDef$defnId,
+           formula=coxDef$formula, localServer=TRUE)
> for (i in seq.int(nSites)) {
+     master$addSite(siteNames[i], siteURLs[[i]],
+                 dataFileName=siteDataFiles[[i]])
+ }
> result <- master$run()
```

At the end of the run, a summary can be printed.

```
> print(master$summary(), digits=5)
        coef exp(coef)  se(coef)       z          p
1 -0.0280495   0.97234 0.0081301 -3.4501 5.6041e-04
2  0.0091441   1.00919 0.0049918  1.8318 6.6979e-02
3 -0.5219296   0.59337 0.1244240 -4.1948 2.7315e-05
4 -0.1941709   0.82352 0.0482507 -4.0242 5.7168e-05
5  0.2636376   1.30166 0.1082448  2.4356 1.4868e-02
6 -0.2400609   0.78658 0.1156319 -2.0761 3.7887e-02
7 -0.2125720   0.80850 0.0937466 -2.2675 2.3359e-02
```

As can be seen, the results are similar to the original model fit.

We have successfully tested other examples in a real three-site configuration, involving two US and one UK site. Tests involving real locally-originated data are underway.

6. Privacy Control and Confidence

The motivating principles guiding the architecture of the distributed computation method we have built include the preservation of data privacy, retention of local control and building confidence in the safety of the method by using an open platform. We list several features that serve that principle.

Registering machines Only known machines are permitted to request computations

Logging All requests for computation are logged by the local computation server, providing accountability. Site specific dashboards can be developed (as **shiny** apps, for example) for analyzing these logs

Single computation server A site could use a single computation server for all distributed collaborations, or a small number, simplifying the oversight task. This simplicity, however, creates a loophole. Our model relies on associating a unique identifier with each *compDef* and associated data at each site. Thus, if slave site A participates in two computations `id1` and `id2` with two different institutions B and C respectively, there is the possibility that B may be able to run a computation on data that was made available to site C and vice-versa if B knows `id2`. However, this is easily handled by means of a generic gatekeeper and reverse proxy application in front of the **opencpu** server and further runtime checks on requests.

Contingent participation A site agrees to or refuses a specific computation, which it can inspect, and can withdraw at any time from any computation, so it retains complete control over the future use of its data.

Limited computation The computation server can only respond to the specific computation request described in the *compdef*. It cannot run other functions or execute other instructions.

7. Some possible use cases

We anticipate that the main initial use will be to make it possible for investigators who are already known to each other to quickly pool information by mounting ”pop-up” collaborations. Once the system is set up, understood and has passed local due diligence for privacy protection, the additional cost and effort required to start a new computation among such a working group should be negligible. As use cases occur, a byproduct will be a growing set of mutually understood data fields, making future computations easier to set up.

There are a few other use cases that we expect to see after some time.

Preparation for aggregation A group of sites who are considering the possibility of aggregating data for some purpose might start their collaboration with a few ”small wins”, to test their ability to rationalize their data dictionaries and see if there is sufficient value in aggregation.

Combining information when control over data use is important Owners of data, such as pharmaceutical companies, who are not willing to transfer them may be willing to participate in collaborations that make use of the distributed technologies. In a single-slave configuration, information from a single source could be used to fit a model proposed by an outside party

(acting as the master). This provides a way for organizations to allow others to access their data (for the purpose of model fitting) in a controlled manner.

Virtual aggregation In a group of sites that often collaborate, a library of "translated terms" and the associated data fields could be developed, creating a virtual data aggregation, whose information would be accessed by distributed model building. If sufficient trust has been created, it is possible to make the setup of a new computation nearly transparent to the investigator.

8. Discussion

The data aggregator must work in a highly restrictive regulatory regime that controls and limits the export of (identified) Protected Health Information (PHI) beyond the originating institution. Owners of data can be reluctant to cede control over the use of their data, even when privacy issues can be resolved. It seems realistic to predict that regulatory concerns, privacy issues, and reluctance to let raw data leave the local source will continue to make it difficult to construct central repositories of data on biomarkers, treatments, and clinical outcomes, despite ongoing efforts by the NIH to support (and even mandate) such repositories. The barriers to access for such data do not seem to be falling. This is of special concern for the aggregators of data (such as gene sequences) that are increasingly likely to emerge as part of the diagnostic and treatment process, and therefore become part of a system of medical records. Such data would be regulated as PHI (under HIPAA) in the future if a consensus develops that they cannot be irreversibly de-identified.

The lack of uniform standards for data management, database ontologies, and the specific content and scope of databases makes global aggregation of complex datsets laborious, expensive, and time-consuming. Reluctance to adopt common standards makes central repositories difficult to achieve, and the effort is seldom funded by outside sources. Investigators might be willing to pay the modest price for standardizing the part of the data that is needed for a particular project, and then to continue to increase common ontological themes as they are rewarded with analyses.

There are limitations to our current implementation. We have not yet built methods beyond the two described above. We are working on others, including the *lasso*. We expect that the repertoire of models will expand as others add their own features and make them available. There is not yet an easy-to-use procedure for handling factor variables whose support varies by site. This and many other improvements in ease of use will require further effort. It is often the case that most data at sites is in databases such as Redcap or i2b2. Our examples have used pure R structures for persisting the data. A simple modification can enable the use of a system such as Redcap. In discussions with personnel at some sites, this configuration seems most acceptable to the CIOs, enabling periodic updates of the data.

We used R for our current implementation since it offers great potential to add new methods. Often, this involves mere re-engineering of already existing algorithms in packages to enable distributed computations. However, a robust implementation of a distributed computation must deal with many modes of failure. More work remains to be done on this, particularly in implementing a robust messaging system to deal with failures.

The current implementation makes uses of the JSON serialization format, simply because it is the best supported format of **opencpu**. It would be better to use a well-established

portable and stable serialization format, such as Google protocol buffers (Google 2012; Eddelbuettel, Stokely, and Ooms 2014; Francois, Eddelbuettel, Stokely, and Ooms 2014). It is straightforward to adapt to this protocol once it is more widely supported in **opencpu** for example.

These limitations of our first implementation notwithstanding, we hope that lowering the barriers to shared statistical computation will accelerate the pace of collaboration and increase the accessibility of information that is now locked up.

9. Acknowledgements

This work by supported in part by grants the Cancer Center Support grant 1 P30 CA124435 from the National Cancer Center, the Clinical and Translational Science Award 1UL1 RR025744 for the Stanford Center for Clinical and Translational Education and Research (Spectrum) from the National Center for Research Resources, National Institutes of Health and award LM07033 from the National Institutes of Health.

References

Anderson E, Bai Z, Bischof C, Blackford S, Demmel J, Dongarra J, Du Croz J, Greenbaum A, Hammerling S, McKenney A, *et al.* (1999). *LAPACK Users' guide*, volume 9. Siam.

Boyd S, Parikh N, Chu E, Peleato B, Eckstein J (2011). "Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers." *Found. Trends Mach. Learn.*, **3**(1), 1–122. ISSN 1935-8237. doi:10.1561/2200000016. URL http://dx.doi.org/10.1561/2200000016.

Chen H, Martin B, Daimon CM, Maudsley S (2013). "Effective use of latent semantic indexing and computational linguistics in biological and biomedical applications." *Frontiers in physiology*, **4**.

Cox DR (1972). "Regression Models and Life-Tables." *Journal of the Royal Statistical Society. Series B (Methodological)*, **34**(2), 187–220. URL http://links.jstor.org/sici?sici=0035-9246%281972%2934%3A2%3C187%3ARMAL%3E2.0.CO%3B2-6.

Eddelbuettel D, Stokely M, Ooms J (2014). "RProtoBuf: Efficient Cross-Language Data Serialization in R." *Arxiv*. doi:arXiv:1401.7372.

Francois R, Eddelbuettel D, Stokely M, Ooms J (2014). *RProtoBuf: R Interface to the Protocol Buffers API. R package version*. doi:arXiv:1406.4806.

Google (2012). *Protocol Buffers: Developer Guide*. URL URLhttp://code.google.com/apis/protocolbuffers/docs/overview.html.

Hosmer D, Lemeshow S, May S (2008). *Applied Survival Analysis*. Wiley. ISBN-13: 0471754994, URL http://www.wiley.com/WileyCDA/WileyTitle/productCd-0471754994.html.

Jiang W, Li P, Wang S, Wu Y, Xue M, Ohno-Machado L, Jiang X (2013). "WebGLORE: a web service for Grid LOgistic REgression." *Bioinformatics (Oxford, England)*, **29**(24), 3238–3240. ISSN 1367-4811. doi:10.1093/bioinformatics/btt559. PMID: 24072732 PMCID: PMC3842761.

Leek JT, Storey JD (2007). "Capturing heterogeneity in gene expression studies by surrogate variable analysis." *PLoS genetics*, **3**(9), e161.

Mazumder R, Hastie T, Tibshirani R (2010). "Spectral regularization algorithms for learning large incomplete matrices." *The Journal of Machine Learning Research*, **11**, 2287–2322.

Murtagh MJ, Demir I, Jenkings KN, Wallace SE, Murtagh B, Boniol M, Bota M, Laflamme P, Boffetta P, Ferretti V, Burton PR (2012). "Securing the data economy: translating privacy and enacting security in the development of DataSHIELD." *Public health genomics*, **15**(5), 243–253. ISSN 1662-8063. doi:10.1159/000336673. PMID: 22722688.

NIH (2014). "NIH Collaboratory: Rethinking Clinical Trials." Accessible at https://www.nihcollaboratory.org.

Novembre J, Johnson T, Bryc K, Kutalik Z, Boyko AR, Auton A, Indap A, King KS, Bergmann S, Nelson MR, *et al.* (2008). "Genes mirror geography within Europe." *Nature*, **456**(7218), 98–101.

Ooms J (2014). "The OpenCPU System: Towards a Universal Interface for Scientific Computing through Separation of Concerns." *Arxiv*. doi:arXiv:1406.4806. URL http://www.opencpu.org.

R Core Team (2014). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL http://www.R-project.org/.

RStudio, Inc (2014). *shiny: Web Application Framework for R*. R package version 0.10.2.2, URL http://CRAN.R-project.org/package=shiny.

Terry M Therneau, Patricia M Grambsch (2000). *Modeling Survival Data: Extending the Cox Model*. Springer, New York. ISBN 0-387-98784-3.

Therneau TM (2014). *A Package for Survival Analysis in S*. R package version 2.37-7, URL http://CRAN.R-project.org/package=survival.

Urbanek S (2013). *Rserve: Binary R server*. R package version 1.7-3, URL http://CRAN.R-project.org/package=Rserve.

Wu Y, Jiang X, Kim J, Ohno-Machado L (2012a). "Grid Binary LOgistic REgression (GLORE): building shared models without sharing data." *Journal of the American Medical Informatics Association: JAMIA*, **19**(5), 758–764. ISSN 1527-974X. doi:10.1136/amiajnl-2012-000862. PMID: 22511014 PMCID: PMC3422844.

Wu Y, Jiang X, Ohno-Machado L (2012b). "Preserving Institutional Privacy in Distributed binary Logistic Regression." *AMIA ... Annual Symposium proceedings / AMIA Symposium. AMIA Symposium*, **2012**, 1450–1458. ISSN 1942-597X. PMID: 23304425 PMCID: PMC3540539.

A. The computational method for a distributed Cox PH fit

We focus on survival data, where the number of covariates might be gathered on subjects that are followed over time. The random quantity of interest is the time to "event" T_i for each subject i which may be observed in some patients (in which case the observation is uncensored) or not (in which case it is censored). Thus, each subject i yields data (T_i, δ_i) where $\delta_i = 1$ if event, 0 otherwise, independent of T_i in addition to covariates that may vary over time.

The Cox proportional hazards model assumes a hazard function of the form

$$\boldsymbol{\lambda}(t) = \lambda_0(t) \exp(\boldsymbol{X}\boldsymbol{\beta}), \tag{1}$$

where $\boldsymbol{\lambda}$ is an $n \times 1$ vector, n being the number of subjects, λ_0, an unspecified baseline hazard, $\boldsymbol{X}$ the $n \times p$ vector of covariates and $\boldsymbol{\beta}$ a $p \times 1$ vector of coefficients.

Model fitting and inference is accomplished by maximizing a partial likelihood function. The explicit form of this likelihood can be written out using a counting process formulation following Therneau and Grambsch (Terry M. Therneau and Patricia M. Grambsch 2000).

For each subject i, let $N_i(t)$ be the number of observed events in $[0, t]$ and $Y_i(t)$, the indicator of whether the subject is in the "risk set", that is, the subject may potentially contribute an event at time t. Then the log of the partial likelihood function introduced by Cox (Cox 1972) has the form:

$$l(\boldsymbol{\beta}|\boldsymbol{X}) = \sum_{i=1}^{n} \int_0^{\infty} \left[Y_i(t)\boldsymbol{X}_i(t)\boldsymbol{\beta} - \log\left(\sum_j Y_j(t) r_j(\boldsymbol{\beta}, t) \right) \right] dN_i(t) \tag{2}$$

where i indexes the subject, $r_i(\boldsymbol{\beta}, t) = \exp[\boldsymbol{X}_i\boldsymbol{\beta}]$ is the risk score, $\boldsymbol{X}_i(t)$ is the covariate row vector. The score vector is given by:

$$U(\boldsymbol{\beta}|\boldsymbol{X}) = \sum_{i=1}^{n} \int_0^{\infty} \boldsymbol{Z}_i(\boldsymbol{\beta}, s) dN_i(s) \tag{3}$$

where

$$\boldsymbol{Z}_i(\boldsymbol{\beta}, s) = \boldsymbol{X}_i(s) - \overline{x}(\boldsymbol{\beta}, s) \tag{4}$$

and

$$\overline{x}(\boldsymbol{\beta}, s) = \frac{\sum Y_i(s) r_i(\boldsymbol{\beta}, s) \boldsymbol{X}_i(s)}{\sum Y_i(s) r_i(\boldsymbol{\beta}, s)}. \tag{5}$$

The negative of the hessian is the Fisher information given by

$$\boldsymbol{I}(\boldsymbol{\beta}|\boldsymbol{X}) = \sum_{i=1}^{n} \int_0^{\infty} V(\boldsymbol{\beta}, s) dN_i(s) \tag{6}$$

where

$$V(\boldsymbol{\beta}, s) = \frac{\sum Y_i(s) r_i(\boldsymbol{\beta}, s) \left[\boldsymbol{Z}_i^T(\boldsymbol{\beta}, s) \boldsymbol{Z}_i(\boldsymbol{\beta}, s) \right]}{\sum Y_i(s) r_i(\boldsymbol{\beta}, s)}. \tag{7}$$

The solution $\hat{\boldsymbol{\beta}}$ to the score equation 3 and the inverse of the information matrix $I^{-1}(\hat{\boldsymbol{\beta}})$ are used as the estimate and the variance respectively along with the fact that the estimate

Data: $\boldsymbol{X} \in \mathcal{R}^{n \times p}$
Result: $u \in \mathcal{R}^n$, $v \in \mathcal{R}^p$, and $d > 0$
$u \leftarrow (\frac{1}{\sqrt{n}}, \frac{1}{\sqrt{n}}, \ldots, \frac{1}{\sqrt{n}})$;
repeat
  $v \leftarrow \boldsymbol{X}^T u$;
  $v \leftarrow v/\|v\|$;
  $u \leftarrow \boldsymbol{X} v$;
  $d \leftarrow \|u\|$;
  $u \leftarrow u/\|u\|$;
until *convergence*;

Algorithm 1: Alternating Algorithm for rank-1 SVD

is asymptotically normally distributed. One can use a Newton-Raphson method to solve equation 3 with an initial value for $\boldsymbol{\beta}$, often just $\mathbf{0}$. The method is quite robust and in fact implemented in the popular `R survival` package (Therneau 2014) of Terry Therneau.

We are interested in *stratified* Cox models where the data is divided into several strata, each with its own baseline hazard, yet use a common set of coefficients $\boldsymbol{\beta}$. Specifically, if there are K strata and one views the entire $n \times p$ matrix $\boldsymbol{X}$ as a stacking of k sub-matrices $\boldsymbol{X}^{[k]}$ each of dimension $n_k \times p$ with $n = \sum_{k=1}^{K} n_k$, then subject i in stratum k incurs a hazard $\lambda_k(t) \exp(\boldsymbol{X}_i^{[k]} \boldsymbol{\beta})$. Such models are particularly applicable in multicenter studies and trials where the one needs to account for patient mix at each institution, for example. In this case, the overall log-likelihood is a sum over the strata

$$l(\boldsymbol{\beta}|\boldsymbol{X}) = \sum_{k=1}^{K} l(\boldsymbol{\beta}, \boldsymbol{X}^{[k]}). \tag{8}$$

where each component of the sum is in fact given by equation 2. It follows that the score function and the Fisher information also partition into sums:

$$U(\boldsymbol{\beta}|\boldsymbol{X}) = \sum_{k=1}^{K} U(\boldsymbol{\beta}, \boldsymbol{X}^{[k]}) \tag{9}$$

and

$$I(\boldsymbol{\beta}|\boldsymbol{X}) = \sum_{k=1}^{K} I(\boldsymbol{\beta}, \boldsymbol{X}^{[k]}) \tag{10}$$

where the components in each sum above are given by equations 3 and 6.

Note that in each stratum k, the computation of the likelihood, score and information for the site-stratified Cox proportional hazards model uses only stratum data $\boldsymbol{X}^{[k]}$. This feature of the site-stratified model enables distributed computation, since each site in a distributed optimization of the partial likelihood only need provide values of $l(\boldsymbol{\beta}, \boldsymbol{X}^{[k]})$, $U(\boldsymbol{\beta}|\boldsymbol{X}^{[k]})$, and $I(\boldsymbol{\beta}|\boldsymbol{X}^{[k]})$ for a master process to estimate $\boldsymbol{\beta}$ and its variance.

B. The computational method for a distributed rank-k SVD

The singular value decomposition (SVD) is a method for summarizing a dataset. The SVD of $\boldsymbol{X}$, the $n \times p$ matrix of covariates, is formed by the matrices $\boldsymbol{U}$, $\boldsymbol{V}$, $\boldsymbol{D}$ such that

$$\boldsymbol{X} = \boldsymbol{U}\boldsymbol{D}\boldsymbol{V}^T,\ \boldsymbol{U}^T\boldsymbol{U} = I,\ \boldsymbol{V}^T\boldsymbol{V} = I,\ \text{and } \boldsymbol{D} \text{ is diagonal with entries.} \tag{11}$$

Data: each slave has private data $\boldsymbol{X}_i \in \mathcal{R}^{n_i \times p}$
Result: $V \in \mathcal{R}^{p \times k}$, and $d_1 \geq \ldots d_k \geq 0$
$V \leftarrow 0$, $d \leftarrow 0$ **foreach** *slave site* j **do**
  $U^{[j]} = 0$;
  transmit n_j to master;
end
for $i \leftarrow 1$ **to** k **do**
  foreach *slave site* j **do** $u^{[j]} \leftarrow (1, 1, \ldots, 1)$ of length n_j $\|u\| \leftarrow \sqrt{\sum_j n_j}$;
  transmit $\|u\|, V$, and D to slaves;
  repeat
    foreach *slave site* j **do**
      $u^{[j]} \leftarrow u^{[j]}/\|u\|$;
      calculate $v^{[j]} \leftarrow (\boldsymbol{X}^{[j]} - U^{[j]}DV^T)^T u^{[j]}$;
      transmit $v^{[j]}$ to master;
    end
    $v \leftarrow \sum_j v^{[j]}$;
    $v \leftarrow v/\|v\|$;
    transmit v to slaves;
    foreach *slave site* j **do**
      calculate $u^{[j]} \leftarrow \boldsymbol{X}^{[j]}v$;
      transmit $\|u^{[j]}\|$ to master;
    end
    $\|u\| \leftarrow \sum_j \|u^{[j]}\|$;
    transmit $\|u\|$ to slaves;
    $d_i \leftarrow \|u\|$;
  until *convergence*;
  $V \leftarrow \texttt{cbind}(V, v)$;
  foreach *slave site* j **do** $U^{[j]} \leftarrow \texttt{cbind}(U^{[j]}, u^{[j]})$
end

Algorithm 2: Privacy Preserving Algorithm for rank-k SVD

Additionally, we assume that the values of $\boldsymbol{D}$ is nonnegative and sorted, namely that $\boldsymbol{D} = \texttt{diag}(d_1, d_2, \ldots, d_{\min(n,p)})$ with $d_1 \geq d_2, \geq \cdots \geq 0$. With these assumptions, the SVD of $\boldsymbol{X}$ is unique up to the signs of the columns of $\boldsymbol{U}$ and $\boldsymbol{V}$.

The SVD is a sort of factor model which decomposes the variance of $\boldsymbol{X}$ into what are called principle components. v_1, the first column of $\boldsymbol{V}$, is called the first principle component of X and is the maximizer of $\text{var}(\boldsymbol{X}v)$. The values of u_1 can be viewed as loading that tell you how much of the factor v_1 is present in each observation. $d_1^2/\sum_j d_j^2$ is the proportion of the variance of $\boldsymbol{X}$ that can be explained by v_1, which means that the SVD also sorts the

components in order of their contribution to variance. When we take a rank-k SVD we only keep the k most important factors, equivalent to setting $d_{k+1} = d_{k+2} = \cdots = d_{\min(n,p)} = 0$.

The SVD has many uses in medical research. Often times, the results are themselves of interest for examining a factor model. This is the case for methods like Latent Semantic Indexing (Chen, Martin, Daimon, and Maudsley 2013). One recent example of an SVD that was done on genomic data is (Novembre, Johnson, Bryc, Kutalik, Boyko, Auton, Indap, King, Bergmann, Nelson *et al.* 2008) where the authors are able to reconstruct a map of europe from the first two principal components of SNP data. Additionally, SVDs are often used as a preprocessing step before some further data analysis. There are many examples of this, including Principal Component Regression, SVD based imputation methods, and Surrogate Variable Analysis – a technique based on trying to remove unwanted variation from high dimensional data (Mazumder, Hastie, and Tibshirani 2010; Leek and Storey 2007).

Calculating an SVD can be a challenging problem from a numerical stability standpoint. For this reason, it is relatively rare for people to implement their own version of an SVD; R and Matlab both use the LAPACK implementation (Anderson, Bai, Bischof, Blackford, Demmel, Dongarra, Du Croz, Greenbaum, Hammerling, McKenney *et al.* 1999). Most of these methods involve transforming the $\boldsymbol{X}$ matrix in a way that cannot be done while maintaining privacy. That said, there is a relatively simple iterative procedure for calculating the rank-1 SVD that can be modified to be privacy preserving. Algorithm 1 calculates (u, v, d), the rank-1 SVD of $\boldsymbol{X}$.

Since singular vectors can be found successively by removing the effect of the top singular vector and then finding the rank-1 approximation again, we can repeatedly apply the algorithm 1 in order to get a rank-k SVD. Combining that idea with some modifications to the previous algorithm in order to preserve privacy, we get algorithm 2 for rank-k privacy protecting SVD.

Affiliation:

Balasubramanian Narasimhan
Department of Health Research and Policy, and
Department of Statistics
Stanford University
Stanford, CA 94305
E-mail: naras@stat.stanford.edu
URL: http://statweb.stanford.edu/~naras/

Daniel L. Rubin
Department of Radiology
Stanford University
Stanford, CA 94305
E-mail: dlrubin@stanford.edu
URL: http://stanford.edu/~rubin/

Samuel M. Gross
Department of Statistics
Stanford University
Stanford, CA 94305

Marina Bendersky
Department of Radiology
Stanford University
Stanford, CA 94305

Philip W. Lavori
Department of Health Research and Policy
Stanford University
Stanford, CA 94305
E-mail: lavori@stanford.edu
URL: https://med.stanford.edu/profiles/philip-lavori